

Optimal Index Policies for Anomaly Localization in Resource-Constrained Cyber Systems

Kobi Cohen¹, Qing Zhao¹, *Fellow, IEEE*, and Ananthram Swami², *Fellow, IEEE*

Abstract

The problem of anomaly localization in a resource-constrained cyber system is considered. Each anomalous component of the system incurs a cost per unit time until its anomaly is identified and fixed. Different anomalous components may incur different costs depending on their criticality to the system. Due to resource constraints, only one component can be probed at each given time. The observations from a probed component are realizations drawn from two different distributions depending on whether the component is normal or anomalous. The objective is a probing strategy that minimizes the total expected cost, incurred by all the components during the detection process, under reliability constraints. We consider both independent and exclusive models. In the former, each component can be abnormal with a certain probability independent of other components. In the latter, one and only one component is abnormal. We develop optimal index policies under both models. The proposed index policies apply to a more general case where a subset (more than one) of the components can be probed simultaneously. The problem under study also finds applications in spectrum scanning in cognitive radio networks and event detection in sensor networks.

Index Terms—Anomaly localization, sequential hypothesis testing, Sequential Probability Ratio Test (SPRT), composite hypothesis testing.

This work was supported by Army Research Laboratory under Grant W911NF1120086 and by the National Science Foundation under Grant CCF-1320065.

Part of this work was presented at the IEEE Global Conference on Signal and Information Processing (GlobalSIP), Austin, Texas, USA, December 2013.

¹ Department of Electrical and Computer Engineering, University of California, Davis, CA 95616 USA (e-mail: yscohen@ucdavis.edu; qzhao@ucdavis.edu).

² Army Research Laboratory, Adelphi, MD 20783 USA (e-mail: a.swami@ieee.org).

I. INTRODUCTION

Consider a cyber system with K components. Each component may be in a normal or an abnormal state. If abnormal, component k incurs a cost c_k per unit time until its anomaly is identified and fixed. Due to resource constraints, only one component can be probed at a time, and switching to a different component is allowed only when the state of the currently probed component is declared. The observations from a probed component (say k) follow distributions $f_k^{(0)}$ or $f_k^{(1)}$ depending on whether the component is normal or anomalous, respectively. The objective is a probing strategy that dynamically determines the order of the sequential tests performed on all the components so that the total cost incurred to the system during the entire detection process is minimized under reliability constraints.

A. Main Results

The above problem presents an interesting twist to the classic sequential hypothesis testing problem. In the case when there is only one component, minimizing the cost is equivalent to minimizing the detection delay, and the problem is reduced to a classic sequential test where both the simple and the composite hypothesis cases have been well studied. With multiple components, however, minimizing the detection delay of each component is no longer sufficient. The key to minimizing the total cost is the order at which the components are being tested. It is intuitive that we should prioritize components that incur higher costs when abnormal as well as components with higher prior probabilities of being abnormal. Another parameter that plays a role in the total system cost is the expected time in detecting the state of a component, which depends on the observation distributions $\{f_k^{(0)}, f_k^{(1)}\}$. It is desirable to place components that require longer testing time toward the end of the testing process. The challenge here is how to balance these parameters in the dynamic probing strategy.

We show in this paper that the optimal probing strategy is an open-loop policy where the testing order can be predetermined, independent of the realizations of each individual test in terms of both the test outcome and the detection time. Furthermore, the probing order is given by a simple index. Specifically, under the independent model where each component is abnormal with a prior probability π_k independent of other components, the index is in the form of $\pi_k c_k / \mathbf{E}(N_k)$, where $\mathbf{E}(N_k)$ is the expected detection time for component k . Under the exclusive model where one and only one component is abnormal, the index is in the form of $\pi_k c_k / \mathbf{E}(N_k | H_0)$ where $\mathbf{E}(N_k | H_0)$ is the expected detection time for component k under the hypothesis of it being normal. These index forms give a clean expression on how the three key parameters—the cost, the prior probability, and the difficulty in distinguish normal distribution $f_k^{(0)}$ from abnormal distribution $f_k^{(1)}$ —are balanced in choosing the probing order. Furthermore, it is interesting to

notice the difference in the indices for these two models. Intuitively speaking, under the independent model, the detection time of any component, normal or abnormal, adds to the delay in catching the next abnormal component, while under the exclusive model, only the detection times of components in a normal state adds to the delay in catching the abnormal component.

The above simple index forms of the probing order are optimal for both the simple hypothesis ($\{f_k^{(0)}, f_k^{(1)}\}_{k=1}^K$ are known) and the composite hypothesis ($\{f_k^{(0)}, f_k^{(1)}\}_{k=1}^K$ have unknown parameters) cases. These index policies also apply to the case where multiple components can be probed simultaneously. While the optimality of these indices in this case remains open, simulation examples demonstrate their strong performance.

B. Applications

The problem considered here finds applications in anomaly detection in cyber systems, spectrum scanning in cognitive radio systems, and event detection in sensor networks. In the following we give two specific examples.

Consider a cyber network consisting of K components (which can be routers, paths, etc.). Due to resource constraints, only a subset of the components can be probed at a time. An Intrusion Detection System (IDS) analyzes the traffic over the components to detect Denial of Service (DoS) attacks (such attacks aim to overwhelm the component with useless traffic to make it unavailable for its intended use). Let the cost c_k be the normal expected traffic (packets per unit time) over component k . Thus, in this example minimizing the total expected cost minimizes the total expected number of failed packets in the network during DoS attacks. The exclusive model applies to cases where an intrusion to a subnet, consisting of K components, has been detected and the probability of each component being compromised is small (thus with high probability, there is only one abnormal component).

Another example is spectrum sensing in cognitive radio systems. Consider a spectrum consisting of K orthogonal channels. Accessing an idle channel leads to a successful transmission, while accessing a busy channel results in a collision with other users. A Cognitive Radio (CR) is an intelligent device that can detect and access idle channels in the wireless spectrum [1]. Due to resource constraints, only a subset of the channels can be sensed at a time. Once a channel is identified as idle, the CR transmits over it. Let c_k be the achievable rate over channel k . Thus, in this example minimizing the total expected cost minimizes the total expected loss in data rate during the spectrum sensing process.

C. Related Work

The classic sequential hypothesis testing problem which pioneered by Wald [2] considers only a single stochastic process. For simple binary hypothesis testing, Wald showed that the Sequential Probability Ratio Test (SPRT) is optimal in terms of minimizing the expected sample size under given type I and type II error probability constraints. Various extensions for M-ary hypothesis testing and composite hypothesis testing were studied in [3]–[9] for a single process. In these cases, asymptotically optimal performance can be obtained as the error probability approaches zero.

A number of studies exist in the literature that consider sequential detection over multiple processes. Differing from this work that focuses on minimizing the total cost incurred by anomalous components, these existing results adopt the objective of minimizing the total detection delay. In particular, the problem of quickly detecting an idle period over multiple independent ON/OFF processes was considered in [10] where a threshold policy was shown to be optimal. The ON/OFF nature of the processes and the objective of minimizing the total detection delay make the problems considered in [10] fundamentally different from the one considered in this work. In [11], the problem of quickest detection of the emergence of primary users in multi-channel cognitive radio networks was considered. In [12], the problem of quickest detection of idle channels over K independent channels was studied. The idle/busy state of each channel was assumed fixed over time, and the objective was to minimize the detection delay under error constraints. It was shown that the optimal policy is to carry out an independent SPRT over each channel, irrespective of the testing order. In contrast to [12], we show in this paper that the optimal policy in our model highly depends on the testing order even when the processes are independent. In [13], the problem of identifying the first abnormal sequence among an infinite number of i.i.d sequences was considered. An optimal cumulative sum (CUSUM) test was established under this setting. Variations of the latter model have been studied in [14], [15]. The sequential search problem under the exclusive model was investigated in [16]–[19]. Optimal policies were derived for the problem of quickest search over Wiener processes [16]–[18]. It was shown in [16], [17] that the optimal policy is to select the sequence with the highest posterior probability of being the target at each given time. In [18], an SPRT-based solution was derived, which is equivalent to the optimal policy in the case of searching over Wiener processes. However, minimizing the total expected cost in our model leads to a different problem and consequently a different index policy.

The classic target whereabouts problem is also a detection problem over multiple processes. In this problem, multiple locations are searched to locate a target. The problem is often considered under the

setting of fixed sample size as in [20]–[23]. In [20], [21], [23], searching in a specific location provides a binary-valued measurement regarding the presence or absence of the target. In [22], Castanon considered the dynamic search problem under continuous observations: the observations from a location without the target and with the target have distributions f and g , respectively. The optimal policy was established under a symmetry assumption that $f(x) = g(b - x)$ for some b .

The anomaly detection problem may also be considered as a variation of active hypothesis testing in which the decision maker chooses and dynamically changes its observation model among a set of observation options. Classic and more recent studies of general active hypothesis testing problems can be found in [24]–[30].

D. Organization

In Section II we describe the system model and problem formulation. In Section III we propose a two-stage optimization problem that simplifies computation while preserving optimality. In Section IV we derive optimal algorithms under the independent and exclusive models for the simple hypothesis case. In Section V we extend our results to the composite hypothesis case: we derive asymptotically optimal algorithms under the independent and exclusive models. In Section VI we provide numerical examples to illustrate the performance of the algorithms.

II. SYSTEM MODEL AND PROBLEM FORMULATION

Consider a cyber system consisting of K components, where each component may be in a normal state or abnormal state. Define

$$\begin{aligned}\mathcal{H}_0 &\triangleq \{k : 1 \leq k \leq K, \text{ component } k \text{ is healthy}\} , \\ \mathcal{H}_1 &\triangleq \{k : 1 \leq k \leq K, \text{ component } k \text{ is abnormal}\} ,\end{aligned}\tag{1}$$

as the sets of the normal and abnormal components.

We consider two different anomaly models.

- 1) Exclusive model: One and only one component is abnormal; the *a priori* probability that component k is the abnormal one is π_k , where $\sum_{k=1}^K \pi_k = 1$.
- 2) Independent model: Each component k is abnormal with *a priori* probability π_k independent of other components.

Every abnormal component k incurs a cost c_k ($0 \leq c_k < \infty$) per unit time until it is tested and identified. Components in a normal state do not incur cost. We focus on the case where only one component can

be probed at a time. The resulting probing strategies apply to the case where a subset of the components can be probed simultaneously and their performance in this case are studied via simulation examples, given in Section VI. When component k is tested at time t , a measurement (or a vector of measurements) $y_k(t)$ is obtained and is independent over time. If component k is healthy, $y_k(t)$ follows distribution $f_k^{(0)}$; if component k is abnormal, $y_k(t)$ follows distribution $f_k^{(1)}$. We focus first on the simple hypothesis case, where the distributions $f_k^{(0)}, f_k^{(1)}$ are known. In Section V we extend our results to the composite hypothesis case, where the distributions have unknown parameters. We consider the case where switching across components is allowed only when the state of the currently probed component is declared.

The detection process starts at time $t = 1$. The random sample size required to make a decision regarding the state of component k is denoted by N_k . We define τ_k as the stopping time (counted from the beginning of the first test at $t = 1$), at which the decision maker stops taking observations from component k and declares its state. The vector of stopping times for the K components is denoted by $\boldsymbol{\tau} = (\tau_1, \dots, \tau_K)$. For example, assume that $K = 3$ and the decision maker tests the components according to the following order: 3, 1, 2. Then, $\tau_3 = N_3$, $\tau_1 = N_3 + N_1$, $\tau_2 = N_3 + N_1 + N_2$.

Let $\delta_k \in \{0, 1\}$ be a decision rule, which the decision maker uses to declare the state of component k at time τ_k . $\delta_k = 0$ if the decision maker declares that component k is in a healthy state (i.e., H_0), and $\delta_k = 1$ if the decision maker declares that component k is in an abnormal state (i.e., H_1). The vector of decision rules for the K components is denoted by $\boldsymbol{\delta} = (\delta_1, \dots, \delta_K)$.

Let $\mathcal{K}(t)$ be the set of components whose states have not been declared by time t and $\phi(t)$ the index of the component being tested at time t (i.e., a selection rule). Let $\mathbf{y}(t) = \{\phi(i), y_{\phi(i)}(i)\}_{i=1}^t$ be the set of all observations and actions up to time t . The selection rule $\phi(t)$ is a mapping from $\mathbf{y}(t-1)$ to $\mathcal{K}(t)$, indicating which component is chosen to be tested at time t among the components whose states have not been determined. Since switching across components is allowed only when the state of the currently probed component is declared, the selection rule satisfies $\phi(\tau_k - t) = \phi(\tau_k)$ for all $1 \leq t \leq N_{\phi(\tau_k)} - 1$, $k = 1, 2, \dots, K$. The vector of selection rules over the time series is denoted by $\boldsymbol{\phi} = (\phi(1), \phi(2), \dots)$. An admissible strategy \mathbf{s} is a sequence of K sequential tests for the K components and denoted by the tuple $\mathbf{s} = (\boldsymbol{\tau}, \boldsymbol{\delta}, \boldsymbol{\phi})$.

The problem is to find a strategy \mathbf{s} that minimizes the total expected cost, incurred by all the abnormal components during the entire detection process, subject to type I (false-alarm) and type II (miss-detect)

error constraints for each component:

$$\begin{aligned} \inf_{\mathbf{s}} \quad & \mathbf{E} \left\{ \sum_{k \in \mathcal{H}_1} c_k \tau_k \right\} \\ \text{s.t.} \quad & P_k^{FA} \leq \alpha_k, \quad P_k^{MD} \leq \beta_k \quad \forall k = 1, \dots, K, \end{aligned} \quad (2)$$

We point out that the total cost defined in (2) does not include the cost incurred by miss-detected abnormal components. Since the error constraints are typically required to be small, (2) well approximates the actual loss in practice.

We have adopted a model where switching across components is allowed only when the test of a currently chosen component is complete. This model is desirable in practical scenarios when switching among components results in additional cost or delay. This model also reduces the memory requirement since only observations from a single component need to be stored. Furthermore, this model is advantageous from a computational complexity perspective. Detection problems involving multiple processes are partially-observed Markov decision processes (POMDP) [22] which have exponential complexity in general. As a result, computing optimal policies is intractable (except for some special observation distributions as considered in [16], [22]). Thus, simplifying the search model is necessary to make the problem mathematically tractable and provide insights and general design guidelines. Similar assumptions have been adopted in [13], [18], [19] to simplify the search model under different objectives.

III. DECOUPLING OF ORDERING AND SEQUENTIAL TESTING

In this section, we show that the probing order and the sequential testing of each component can be decoupled. As a consequence, the solution to (2) can be obtained in two stages.

In the first stage, we solve the following optimization problem for every component k :

$$\begin{aligned} \inf_{N_k, \delta_k} \quad & \mathbf{E}(N_k | H_i), \quad i = 0, 1 \\ \text{s.t.} \quad & P_k^{FA} \leq \alpha_k, \quad P_k^{MD} \leq \beta_k. \end{aligned} \quad (3)$$

In the second stage, the problem is to find a selection rule ϕ that minimizes the objective function, given the solution to the K subproblems specified in (3):

$$\inf_{\phi} \quad \mathbf{E} \left\{ \sum_{k \in \mathcal{H}_1} c_k \tau_k \mid (\mathbf{N}^*, \boldsymbol{\delta}^*) \right\} \quad (4)$$

where

$$\mathbf{N}^* = (N_1^*, \dots, N_K^*), \quad \boldsymbol{\delta}^* = (\delta_1^*, \dots, \delta_K^*) \quad (5)$$

denote the vectors of stopping times and decision rules, respectively, that solve the K subproblems given in (3). Note that the stopping times $\tau = (\tau_1, \dots, \tau_K)$ are completely specified by N^* and the selection rule ϕ^* that solves (4).

The formulation of the two-stage optimization problem allows us to decompose the original optimization problem (2) into $K + 1$ subproblems (3) and (4). In subsequent sections we show that the two-stage optimization problem preserves optimality under both the independent and exclusive models.

IV. THE SIMPLE HYPOTHESIS CASE

In this section we derive optimal solutions to both the independent and exclusive models when the observation distributions under both hypotheses are completely known. We discuss the implementation of the optimal policies in Section IV-C.

A. SPRT for Each Component

For the simple hypothesis case, the solution to the first stage optimization problem (3) is given by the SPRT [2] as follows.

Assume that the state of component j has been declared at time τ_j and component k is chosen to be tested at time $\tau_j + 1$. Let

$$L_k(n) = \frac{\prod_{t=\tau_j+1}^{\tau_j+n} f_k^{(1)}(y_k(t))}{\prod_{t=\tau_j+1}^{\tau_j+n} f_k^{(0)}(y_k(t))} \quad (6)$$

be the Likelihood Ratio (LR) between the two hypotheses for component k at stage n .

In SPRT, the stopping and decision rules are given by comparing the LR with boundary values at each stage n [2]. Specifically, let A_k, B_k ($B_k > 1/A_k$) be the boundary values used by the SPRT for component k . The SPRT algorithm is carried out as follows:

- If $L_k(n) \in ((A_k)^{-1}, B_k)$, continue to take observations from component k .
- If $L_k(n) \geq B_k$, stop taking observations from component k and declare it as abnormal (i.e., $\delta_k = 1$).

Clearly, $N_k = n$.

- If $L_k(n) \leq (A_k)^{-1}$, stop taking observations from component k and declare it as normal (i.e., $\delta_k = 0$). Clearly, $N_k = n$.

Implementation of the SPRT requires the computation of A_k and B_k to ensure the constraints on the error probabilities. In general, the exact determination of the boundary values is laborious and depends on the observation distribution. Wald's approximation can be applied to simplify the computation [2]:

$$B_k \approx \frac{1 - \beta_k}{\alpha_k} \quad , \quad A_k \approx \frac{1 - \alpha_k}{\beta_k} . \quad (7)$$

Wald's approximation performs well for small α_k, β_k and is asymptotically optimal as α_k, β_k approach zero. Since type *I* and type *II* errors are typically required to be small, Wald's approximation is widely used in practice [2].

B. Optimal Index Policies

We now consider the second stage optimization problem specified by (4) and (5). Our main result is to establish the optimal selection rule as the πcN -rule for the independent model and the πcN_0 rule for the exclusive model. Specifically, the πcN -rule dictates that the components be tested in a decreasing order of $\pi_k c_k / \mathbf{E}(N_k)$ and the πcN_0 -rule dictates that the components be tested in a decreasing order of $\pi_k c_k / \mathbf{E}(N_k | H_0)$. Note that these optimal selection rules are open loop policies: the testing orders can be determined offline (see Section IV-C for the computation of the indices). With the optimal solution to (3), the optimal anomaly detection strategy is to carry out a series of SPRTs on the components ordered according to either the πcN -rule or the πcN_0 -rule. The resulting strategies are thus referred to as πcN -SPRT and πcN_0 -SPRT.

The index selection rules πcN and πcN_0 are intuitively satisfying. The priority of component k in terms of testing order should be higher as the cost c_k increases, or the *a priori* probability of being abnormal π_k increases. Under the independent model, the priority of component k in terms of testing order should be higher as the expected sample size $\mathbf{E}(N_k)$ decreases (since $\mathbf{E}(N_k)$ contributes to the cost of every component which is tested after component k). On the other hand, under the exclusive model, the priority of component k in terms of testing order depends on $\mathbf{E}(N_k | H_0)$ rather than $\mathbf{E}(N_k)$. The reason is that if component k is abnormal, there is no additional cost, incurred by other components (since only one component is abnormal). On the other hand, if component k is healthy, then $\mathbf{E}(N_k | H_0)$ contributes to the cost of the components (which may be abnormal) tested after component k .

The optimality of the algorithms is shown in the following theorem.

Theorem 1: Under the independent and exclusive models, the πcN -SPRT and πcN_0 -SPRT algorithms, respectively, solve the original optimization problem (2).

Proof: See Appendices VIII-A and VIII-B. ■

While πcN -rule and πcN_0 -rule are open-loop policies, Theorem 1 shows that they are optimal among the class of both open-loop and closed-loop selection rules. It should be noted that open-loop policies may not preserve optimality under non-linear cost functions or other correlated models. In these cases, the

optimal testing order might need to be updated dynamically based on the realizations of each individual test in terms of the test outcome or the detection time.

The πcN -rule and πcN_0 -rule bear some similarity with the result developed in [31]. In [31], the problem of ordering independent operations with given processing times was considered. It was shown that the optimal selection rule for the problem of minimizing an expected weighted sum of completion times of all the operations is to select the components in decreasing order of $c_k/\mathbf{E}(N_k)$, where c_k and $\mathbf{E}(N_k)$ are the weight and the expected processing time for operation k , respectively. However, the problem in (4) is different. First, each component may be normal or abnormal (rather than a given processing time with a fixed distribution) and the expected sample size depends on the component state. Second, the objective is to minimize an expected weighted sum of stopping times of abnormal components only. Third, under the exclusive model, the states of the K components are *dependent*. Furthermore, the original optimization (2) also includes the stopping rules which control the expected sample size.

C. Computing the Indices

Arranging the components according to πcN -rule or πcN_0 -rule can be done in $O(K \log K)$ time via sorting algorithms. However, computing the expected sample size $\mathbf{E}(N_k|H_i)$ for all $k = 1, 2, \dots, K$ can be involved. In general, it is difficult to obtain a closed-form expression for $\mathbf{E}(N_k|H_i)$. One way to evaluate $\mathbf{E}(N_k|H_i)$ is to perform off-line simulations (i.e., carrying out K independent SPRTs for the K components). Another way to evaluate $\mathbf{E}(N_k|H_i)$ is to use a closed-form approximation as follows. Since the solution to (3) is given by the SPRT, Wald's approximation can be applied [2]. For every $i, j = 0, 1$, let

$$D_k(i||j) = \mathbf{E}_i \left(\log \frac{f_k^{(i)}(y_k(1))}{f_k^{(j)}(y_k(1))} \right) \quad (8)$$

be the Kullback-Leibler (KL) divergence between the hypotheses H_i and H_j , where the expectation is taken with respect to $f_k^{(i)}$.

The expected sample size conditioned on each hypothesis is well approximated by [2]:

$$\begin{aligned} \mathbf{E}(N_k|H_0) &\approx \frac{(1 - \alpha_k) \log \tilde{A}_k - \alpha_k \log \tilde{B}_k}{D_k(0||1)}, \\ \mathbf{E}(N_k|H_1) &\approx \frac{(1 - \beta_k) \log \tilde{B}_k - \beta_k \log \tilde{A}_k}{D_k(1||0)}, \end{aligned} \quad (9)$$

where $\tilde{A}_k = (1 - \alpha_k)/\beta_k$, $\tilde{B}_k = (1 - \beta_k)/\alpha_k$ are the approximations to A_k, B_k , given in (7). Note that (9) approach the exact expected sample sizes $\mathbf{E}(N_k|H_0) \rightarrow -\log \beta_k / D_k(0||1)$, $\mathbf{E}(N_k|H_1) \rightarrow$

$-\log \alpha_k / D_k(1||0)$ as the error constraints approach zero.

The expected sample size required to make a decision regarding the state of component k is given by:

$$\mathbf{E}(N_k) = \pi_k \mathbf{E}(N_k|H_1) + (1 - \pi_k) \mathbf{E}(N_k|H_0), \quad (10)$$

where the approximation approaches the exact expected sample size for small α_k, β_k .

Note that optimality of the algorithms is preserved as long as the *order* of the indices is preserved (i.e., the exact index values are not required for optimality). Therefore, optimality can be achieved in practice even when Wald's approximation is used.

V. THE COMPOSITE HYPOTHESIS CASE

In the previous section we focused on the simple hypothesis case, where the distribution under both hypotheses are completely known. For this case, the SPRT was applied to solve (3). However, in numerous cases there is uncertainty in the observation distributions.

For example, Consider a one-parameter distribution $f(y|\theta_k)$, where it is required to test $\theta_k < \theta_k^{(0)}$ against $\theta_k > \theta_k^{(1)} > \theta_k^{(0)}$. As discussed in [2], the SPRT can be applied to this problem by testing $\theta_k = \theta_k^{(0)}$ against $\theta_k = \theta_k^{(1)}$, where the boundary values are set such that the error constraints are satisfied at $\theta_k^{(0)}, \theta_k^{(1)}$. For some important cases, such as an exponential family of distributions, this sequential test has the property that type *I* and type *II* errors are less than α_k, β_k for all $\theta_k < \theta_k^{(0)}$ and $\theta_k > \theta_k^{(1)}$, respectively. However, while the SPRT minimizes the expected sample size at $\theta_k = \theta_k^{(0)}, \theta_k^{(1)}$, it is highly sub-optimal for other values of θ , as demonstrated in Section VI. Therefore, other techniques should be considered under the composite hypothesis case.

Let $\boldsymbol{\theta}_k$ be a vector of unknown parameters of component k . The observations $\{y_k(i)\}_{i \geq 1}$ are drawn from a common distribution $f(y|\boldsymbol{\theta}_k)$, $\boldsymbol{\theta}_k \in \Theta_k$, where Θ_k is the parameter space of component k . If component k is healthy, then $\boldsymbol{\theta}_k \in \Theta_k^{(0)}$; if component k is abnormal, then $\boldsymbol{\theta}_k \in (\Theta \setminus \Theta_k^{(0)})$. Let $\Theta_k^{(0)}, \Theta_k^{(1)}$ be disjoint subsets of Θ_k , where $I_k = \Theta \setminus (\Theta_k^{(0)} \cup \Theta_k^{(1)}) \neq \emptyset$ is an indifference region¹. When $\boldsymbol{\theta}_k \in I_k$, the detector is indifferent regarding the state of component k . Hence, there are no constraints on the error probabilities for all $\boldsymbol{\theta}_k \in I_k$. The hypothesis test regarding component k is to test

$$\boldsymbol{\theta}_k \in \Theta_k^{(0)} \quad \text{against} \quad \boldsymbol{\theta}_k \in \Theta_k^{(1)}.$$

¹The assumption of an indifference region is widely used in the theory of sequential testing of composite hypotheses to derive asymptotically optimal performance. Nevertheless, in some cases this assumption can be removed. For more details, the reader is referred to [5].

Narrowing I_k has the price of increasing the sample size.

Let

$$\begin{aligned}\hat{\theta}_k(n) &= \arg \max_{\theta_k \in \Theta_k} f(\mathbf{y}_k(n) | \theta_k), \\ \hat{\theta}_k^{(i)}(n) &= \arg \max_{\theta_k \in \Theta_k^{(i)}} f(\mathbf{y}_k(n) | \theta_k),\end{aligned}\tag{11}$$

be the Maximum-Likelihood Estimates (MLEs) of the parameters over the parameter spaces Θ_k , $\Theta_k^{(i)}$ at stage n , respectively.

In contrast to the SPRT (for the simple hypothesis case), the theory of sequential tests of composite hypotheses does not provide optimal performance in terms of minimizing the expected sample size under given error constraints. Nevertheless, asymptotically optimal performance can be obtained as the error probability approaches zero.

First, we provide an overview of existing sequential tests for composite hypotheses which are relevant to our problem. Next, we apply these techniques to solve (2).

A. Existing Sequential Tests for Composite Hypothesis Testing

The key idea is to use the estimated parameters to perform a one-sided sequential test to reject H_0 and a one-sided sequential test to reject H_1 . Note that these techniques were introduced for a single process. However, in this paper we apply sequential tests for K components. Thus, we use the subscript k to denote the component index.

1) *Sequential Generalized Likelihood Ratio Test (SGLRT)*: We refer to sequential tests that use the Generalized Likelihood Ratio (GLR) statistics as the SGLRT.

For $i = 0, 1$, let

$$L_k^{(i),GLR}(n) = \log \frac{\prod_{r=1}^n f(y_k(r) | \hat{\theta}_k(n))}{\prod_{r=1}^n f(y_k(r) | \hat{\theta}_k^{(i)}(n))}\tag{12}$$

be the GLR statistics used to reject hypothesis H_i at stage n .

Let

$$N_k^{(i)} = \inf \left\{ n : L_k^{(i),GLR}(n) \geq B_k^{(i)} \right\},\tag{13}$$

be the stopping rule used to reject hypothesis H_i . $B_k^{(i)}$ is the boundary value.

For each component k , the decision maker stops the sampling when $N_k = \min \{N_k^{(0)}, N_k^{(1)}\}$. If $N_k = N_k^{(0)}$, component k is declared as abnormal (i.e., H_0 is rejected). If $N_k = N_k^{(1)}$, component k is declared as normal (i.e., H_0 is accepted).

The SGLRT was first studied by Schwartz [3] for a one-parameter exponential family, who assigned a cost of c for each observation and a loss function for wrong decisions. It was shown that setting

$B_k^{(i)} = \log(c^{-1})$ asymptotically minimizes the Bayes risk as c approaches zero. A refinement was studied by Lai [5], [7], who set a time-varying boundary value $B_k^{(i)} \sim \log((nc)^{-1})$. Lai showed that for a multivariate exponential family this scheme asymptotically minimizes both the Bayes risk and the expected sample size subject to error constraints as c approaches zero [7].

2) *Sequential Adaptive Likelihood Ratio Test (SALRT)*: We refer to sequential tests that use the Adaptive Likelihood Ratio (ALR) statistics as the SALRT.

For $i = 0, 1$, let

$$L_k^{(i), ALR}(n) = \log \frac{\prod_{r=1}^n f(y_k(r) | \hat{\theta}_k(r-1))}{\prod_{r=1}^n f(y_k(r) | \hat{\theta}_k^{(i)}(n))} \quad (14)$$

be the ALR statistics used to reject hypothesis H_i at stage n .

Let

$$N_k^{(i)} = \inf \left\{ n : L_k^{(i), ALR}(n) \geq B_k^{(i)} \right\}, \quad (15)$$

be the stopping rule used to reject hypothesis H_i , where $B_k^{(i)}$ is the boundary value.

For each component k , the decision maker stops the sampling when $N_k = \min \{N_k^{(0)}, N_k^{(1)}\}$. If $N_k = N_k^{(0)}$, component k is declared as abnormal. If $N_k = N_k^{(1)}$, component k is declared as normal.

The SALRT was first introduced by Robbins and Siegmund [4] to design power-one sequential tests. Pavlov used it to design asymptotically (as the error probability approaches zero) optimal (in terms of minimizing the expected sample size subject to error constraints) tests for composite hypothesis testing of the multivariate exponential family [6]. Tartakovsky established asymptotically optimal performance for a more general multivariate family of distributions [8].

The advantage of using the SALRT is that setting $B_k^{(0)} = \log \frac{1}{\alpha_k}$, $B_k^{(1)} = \log \frac{1}{\beta_k}$ satisfies the error probability constraints in (3). However, such a simple setting cannot be applied to the SGLRT. Thus, implementing the SALRT is much simpler than implementing the SGLRT. The disadvantage of using the SALRT is that poor early estimates (for small number of observations) can never be revised even though one has a large number of observations.

B. Asymptotically Optimal Index Policies

It is intuitive that the selection rules in the composite hypothesis case remain the same as in the simple hypothesis case. The resulting strategies are thus referred to as πcN -SGLRT/SALRT and πcN_0 -SGLRT/SALRT algorithms. In the following theorems, we show that the πcN -SGLRT/SALRT and πcN_0 -SGLRT/SALRT algorithms are asymptotically optimal in terms of minimizing the objective function

subject to the error constraints (2) as the error probabilities approach zero². When deriving asymptotics we assume that $P_k^{FA} \rightarrow 0, P_k^{MD} \rightarrow 0$ for all k such that the asymptotic optimality property in terms of minimizing the expected sample size subject to the error constraints holds for each single process for both SGLRT and SALRT, as discussed in Section V-A.

Theorem 2: Consider the independent model under the composite hypothesis case. Let $(\tau^{OPT}, \delta^{OPT}, \phi^{OPT})$ be the optimal solution to (2). Let $(\tau^*, \delta^*, \phi^*)$ be the solution achieved by the πcN -SGLRT/SALRT algorithm. Then, as $P_k^{FA} \rightarrow 0, P_k^{MD} \rightarrow 0$ for all k , we obtain:

$$\begin{aligned} \mathbf{E} \left\{ \sum_{k \in \mathcal{H}_1} c_k \tau_k | (\tau^*, \delta^*, \phi^*) \right\} \\ \sim \mathbf{E} \left\{ \sum_{k \in \mathcal{H}_1} c_k \tau_k | (\tau^{OPT}, \delta^{OPT}, \phi^{OPT}) \right\} \end{aligned} \quad (16)$$

Proof: See Appendix VIII-C. ■

Theorem 3: Consider the exclusive model under the composite hypothesis case. Let $(\tau^{OPT}, \delta^{OPT}, \phi^{OPT})$ be the optimal solution to (2). Let $(\tau^*, \delta^*, \phi^*)$ be the solution achieved by the πcN_0 -SGLRT/SALRT algorithm. Then, as $P_k^{FA} \rightarrow 0, P_k^{MD} \rightarrow 0$ for all k , we obtain:

$$\begin{aligned} \mathbf{E} \left\{ \sum_{k \in \mathcal{H}_1} c_k \tau_k | (\tau^*, \delta^*, \phi^*) \right\} \\ \sim \mathbf{E} \left\{ \sum_{k \in \mathcal{H}_1} c_k \tau_k | (\tau^{OPT}, \delta^{OPT}, \phi^{OPT}) \right\} \end{aligned} \quad (17)$$

Proof: See Appendix VIII-D. ■

C. Computing the Indices

Arranging the components in decreasing order of $\pi_k c_k / \mathbf{E}(N_k)$ or $\pi_k c_k / \mathbf{E}(N_k | H_0)$ requires one to compute the expected sample size $\mathbf{E}(N_k | H_i)$ for all $k = 1, 2, \dots, K$. In general, it is difficult to obtain a closed-form expression for the exact value of $\mathbf{E}(N_k | H_i)$. However, we can use the asymptotic property of the tests to obtain a closed-form approximation of $\mathbf{E}(N_k | H_i)$, which approaches the exact expected sample size as the error probability approaches zero.

²As shown in the proof of Theorems 2, 3, the index policies are still optimal in terms of testing order. The asymptotic optimality is due to the performance of the sequential test under the composite hypothesis case.

For every $i = 0, 1$, let

$$D_k(\boldsymbol{\theta}_k || \boldsymbol{\lambda}) = \mathbf{E}_{\boldsymbol{\theta}_k} \left(\log \frac{f(y_k(1) | \boldsymbol{\theta}_k)}{f(y_k(1) | \boldsymbol{\lambda})} \right) \quad (18)$$

be the KL divergence between the real value of $\boldsymbol{\theta}_k$ and $\boldsymbol{\lambda}$, where the expectation is taken with respect to $f(y | \boldsymbol{\theta}_k)$,

and let

$$D_k^*(\boldsymbol{\theta}_k || \Theta_k^{(i)}) = \inf_{\boldsymbol{\lambda} \in \Theta_k^{(i)}} D_k(\boldsymbol{\theta}_k || \boldsymbol{\lambda}) . \quad (19)$$

Let $I_k^{(0)}, I_k^{(1)}$ be disjoint subsets of I_k and $I_k = I_k^{(0)} \cup I_k^{(1)}$, such that for all $\boldsymbol{\theta}_k \in I_k^{(i)}$ we have $\frac{B_k^{(j)}}{D_k^*(\boldsymbol{\theta}_k || \Theta_k^{(j)})} \leq \frac{B_k^{(i)}}{D_k^*(\boldsymbol{\theta}_k || \Theta_k^{(i)})}$ for $i, j = 0, 1$. Let $P^{(i)}(\boldsymbol{\theta}_k)$ be a prior distribution on $\boldsymbol{\theta}_k$ over $\Theta_k^{(i)} \cup I_k^{(i)}$ (corresponding to H_i). Then, as $P_k^{FA} \rightarrow 0, P_k^{MD} \rightarrow 0$, the conditional expected sample size is given by [6], [7]:

$$\begin{aligned} \mathbf{E}(N_k | H_0) &\sim \int_{\boldsymbol{\theta}_k \in \Theta_k^{(0)} \cup I_k^{(0)}} \frac{B_k^{(1)}}{D_k^*(\boldsymbol{\theta}_k || \Theta_k^{(1)})} dP^{(0)}(\boldsymbol{\theta}_k) , \\ \mathbf{E}(N_k | H_1) &\sim \int_{\boldsymbol{\theta}_k \in \Theta_k^{(1)} \cup I_k^{(1)}} \frac{B_k^{(0)}}{D_k^*(\boldsymbol{\theta}_k || \Theta_k^{(0)})} dP^{(1)}(\boldsymbol{\theta}_k) . \end{aligned} \quad (20)$$

The expected sample size required to make a decision regarding the state of component k is given by:

$$\mathbf{E}(N_k) = \pi_k \mathbf{E}(N_k | H_1) + (1 - \pi_k) \mathbf{E}(N_k | H_0) , \quad (21)$$

which can be well approximated for small error probability using (20).

VI. NUMERICAL EXAMPLES

In this section we present numerical examples to illustrate the performance of the algorithms. Consider a cyber network consisting of K components (which can be routers, paths, etc.), as discussed in Section I-B. Assume that an intruder tries to launch a DoS or Reduction of Quality (RoQ) attacks by sending a large number of packets to a component. RoQ attacks inflict damage on the component, while keeping a low profile to avoid detection. RoQ attacks do not cause denial of service.

To detect such attacks, the IDS performs a traffic-based anomaly detection. It monitors the traffic at each component to decide whether a component is compromised. Roughly speaking, if the actual arrival rate is significantly higher than the arrival rate under the normal state, then the IDS should declare that the component is in an abnormal state. A similar traffic-based detection technique was proposed in [32] for a different model, considering a single process without switching to other components. For each component k , we assume that packets arrive according to a Poisson process with rate $\theta^{(k)}$. When component k is tested, the IDS collects an observation $y_k(n) \in \mathbb{N}_0$ every time unit, which represents the

number of packets that arrived in the interval $(n-1, n)$. Assume that the IDS considers component k as normal if $\theta_k \leq \theta_k^{(0)}$, and tests $\theta_k \leq \theta_k^{(0)}$ against $\theta_k \geq \theta_k^{(1)}$ (i.e., $I_k = \{\theta_k | \theta_k^{(0)} < \theta_k < \theta_k^{(1)}\}$ is the indifference region). We set $c_k = \theta_k^{(0)}$. As discussed in Section I-B, under this setting the optimization problem minimizes the maximal damage to the network in terms of packet-loss.

A. Simple Hypothesis Case

We consider the case where the observations follow Poisson distributions $y_k(n) \sim \text{Poi}(\theta_k^{(0)})$ or $y_k(n) \sim \text{Poi}(\theta_k^{(1)})$ depending on whether component k is healthy or abnormal, respectively, where $\theta_k^{(0)}, \theta_k^{(1)}$ are known to the IDS. To implement the πcN -SPRT and πcN_0 -SPRT algorithms (which are optimal in this scenario for the independent and exclusive models, respectively), we need to compute the LR between the hypotheses, defined in (6), and the expected sample sizes under the hypotheses, which can be well approximated by (9). Let $\Lambda_k(n) = \log L_k(n)$ be the Log-Likelihood Ratio (LLR) between the two hypotheses of component k at stage n , where $L_k(n)$ is defined in (6). After algebraic manipulations, it can be verified that the LLR is given by:

$$\Lambda_k(n) = -n \left(\theta_k^{(1)} - \theta_k^{(0)} \right) + \log \left(\theta_k^{(1)} / \theta_k^{(0)} \right) \sum_{i=1}^n y_k(i). \quad (22)$$

It can be verified that the KL divergence between the hypotheses H_i and H_j , defined in (8), is given by:

$$D_k(i||j) = \theta_k^{(j)} - \theta_k^{(i)} + \theta_k^{(i)} \log \left(\theta_k^{(i)} / \theta_k^{(j)} \right). \quad (23)$$

Substituting (23) in (9) yields the required approximation to the expected sample size. We note that the optimal indices order was preserved using the approximation in (9) under all numerical examples in this section.

Next, we provide numerical examples to illustrate the performance of the algorithms. We compared three schemes: a Random selection SPRT (R-SPRT), where a series of SPRTs are performed until all the components are tested in a random order (which is optimal for the problem of minimizing the detection delay over independent processes [12]), and the proposed πcN -SPRT and πcN_0 -SPRT algorithms, which are optimal under the independent and exclusive models, respectively.

Let $\Delta_K = (100-10)/(K-1)$. We set $c_k = \theta_k^{(0)} = 10 + (k-1)\Delta_K$ (i.e., the costs are equally spaced in the interval $[10, 100]$) and $\theta_k^{(1)} = 1.5 \cdot \theta_k^{(0)}$. The error constraints were set to $P_k^{FA} = 10^{-2}$, $P_k^{MD} = 10^{-6}$ for all k . For the independent and exclusive models, we set $\pi_k = 0.8$ and $\pi_k = 1/K$ for all k , respectively. The performance of the πcN -SPRT and πcN_0 -SPRT algorithms are presented in Fig. 1(a) and 1(b) under the independent and exclusive models, respectively, and compared to the R-SPRT. It can be seen that the

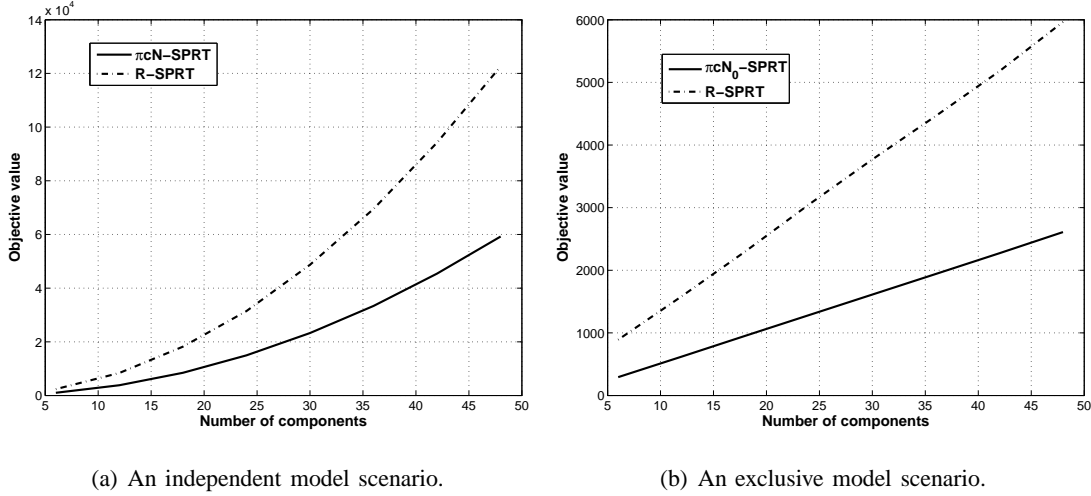


Fig. 1. Objective value as a function of the number of components under the independent and exclusive models.

proposed algorithms save roughly 50% of the objective value as compared to the R-SPRT under both the independent and exclusive model scenarios.

Next, we simulate the independent model when 2 components are observed at a time and the total number of components is $K = 6$. Note that in this case the πcN -SPRT algorithm may not be optimal. We use an exhaustive search as a bench mark to demonstrate the performance of the πcN -SPRT algorithm in this scenario. The exhaustive search is done by performing a sequence of K SPRTs among all the possible testing orders. Then, the minimal objective value is chosen as a bench mark. We set the maximal cost to $c_{max} = 100$ and the costs are equally spaced in the interval $[c_{min}, 100]$. The error constraints were set to $P_k^{FA} = P_k^{MD} = 10^{-2}$ for all k . The performance gain of the exhaustive search scheme over the πcN -SPRT algorithm as a function of c_{min} are presented in Fig. 2. It can be seen that the πcN -SPRT algorithm almost achieves the performance of the exhaustive search scheme in this scenario for all c_{min} . For small c_{min} both algorithms perform the same, since the difference between the indices increases. The exhaustive search outperforms the πcN -SPRT algorithm for $c_{min} > 97$, but the gain remains very small.

B. Composite Hypothesis Case

We consider the case of composite hypotheses, where there is uncertainty in the distribution parameters, as discussed in Section V. To implement the asymptotically optimal the πcN -SGLRT/SALRT and πcN_0 -SGLRT/SALRT algorithms, we need to compute the GLR or ALR statistics, defined in (12), (14) and the




Fig. 2. Performance gain of an exhaustive search over the πcN -SPRT algorithm as a function of c_{\min} under the independent model.

expected sample sizes under the hypotheses, which can be well approximated by (20). The MLEs of the parameters over the parameter spaces Θ_k , $\Theta_k^{(i)}$ are given by the sample mean and the boundary of the alternative parameter space, respectively. As a result, substituting: $\hat{\theta}_k(n) = \frac{1}{n} \sum_{i=1}^n y_k(i)$, $\hat{\theta}_k^{(i)}(n) = \theta_k^{(i)}$, in (12), (14) yields the GLR and ALR statistics, respectively. The KL divergence between the real value of θ_k and the parameter space $\Theta_k^{(i)}$ is given by:

$$D_k^*(\theta_k || \Theta_k^{(i)}) = \theta_k^{(i)} - \theta_k + \theta_k \log \left(\theta_k / \theta_k^{(i)} \right). \quad (24)$$

Substituting (24) in (20) yields the approximate expected sample size.

Next, we provide numerical examples to illustrate the performance of the algorithms under uncertainty. We simulated a network with homogenous components (i.e., any selection rule is optimal). We compared three schemes: R-SPRT, and the πcN -SGLRT/SALRT or πcN_0 -SGLRT/SALRT algorithms (which achieve the same performance in this case) using the SALRT and the SGLRT, discussed in section V-A. We set $\theta_k^{(0)} = 19$, $\theta_k^{(1)} = 21$. Under uncertainty, the IDS considers component k as normal if $\theta_k \leq \theta_k^{(0)}$, and tests $\theta_k \leq \theta_k^{(0)}$ against $\theta_k \geq \theta_k^{(1)}$ (i.e., $I_k = \{\theta_k | 19 < \theta_k < 21\}$ is the indifference region). To implement the SGLRT, we set the cost per observation $c = 10^{-3}$. According to the assigned cost, we obtained the following error probability constraints for all k : $P_k^{FA} \leq 0.026$ for all $\theta^{(k)} \leq 19$ and $P_k^{MD} \leq 0.03$ for all $\theta^{(k)} \geq 21$. We do not restrict the detector's performance for $19 < \theta^{(k)} < 21$ (Note that narrowing the indifference region has the price of increasing the required sample size). In Fig. 3 we show the average number of observations (in a log scale) required for the anomaly detection as a function of $\theta^{(k)}$. As expected, for $\theta_k = 19$ and $\theta_k = 21$ the R-SPRT requires lower sample size

as compared to the proposed schemes. On the other hand, it can be seen that for most values of θ the SGLRT and the SALRT require lower sample size as compared to the R-SPRT. The SALRT performs the worst for $18 < \theta_k < 22$, and performs the best for $\theta_k \notin (18, 22)$, roughly. The SGLRT obtains the best average performance. It can be seen that for large values of θ_k the anomaly is detected very quickly, since the distance between the hypotheses increases. This result confirms that DoS attacks are much easier to detect than RoQ attacks.




Fig. 3. Average number of observations as a function of the arrival rate of packets (denoted by θ).

VII. CONCLUSION

The problem of anomaly localization in a resource-constrained cyber system was studied. Due to resource constraints, only one component can be probed at a time. The observations are realizations drawn from two different distributions depending on whether the component is normal or anomalous. An abnormal component incurs a cost per unit time until it is tested and identified. The problem was formulated as a constrained optimization problem. The objective is to minimize the total expected cost subject to error probability constraints. We considered two different anomaly models: the independent model in which each component can be abnormal independent of other components, and the exclusive model in which there is one and only one abnormal component. For the simple hypothesis case, we derived optimal algorithms for both independent and exclusive models. For the composite hypothesis case, we derived asymptotically (as the error probability approaches zero) optimal algorithms for both independent and exclusive models. These optimal algorithms have low-complexity.

The algorithms developed in this paper can be applied to other models of anomaly detection as well. We can modify the proposed algorithms to any detection scheme that performs a series of tests according

to the πcN -rule or πcN_0 -rule. The required modification is to replace the SPRT/SALRT/SGLRT by any given test. Such modified algorithms minimize the objective function among all the algorithms that perform the given test.

Deriving optimal policies for the anomaly localization problem considered in this paper requires the assumption that switching to a different component is allowed only when the state of the currently probed component is declared. A future research direction is to examine the anomaly localization problem under the case where switching to a different component and declarations of the states of individual components are allowed at all times.

VIII. APPENDIX

In this appendix we provide the proofs for Theorems 1 – 3. For convenience, we use the superscripts $A1, A2$ when referring to the πcN -SPRT and πcN_0 -SPRT algorithms, respectively. We use the superscripts $A3, A4$ when referring to the πcN -SGLRT/SALRT and πcN_0 -SGLRT/SALRT algorithms, respectively.

Throughout the proofs, we use the specific formula for the updated posterior probability of component k being abnormal.

Let $\mathbf{1}_k(n)$ be the probing indicator function, where $\mathbf{1}_k(n) = 1$ if component k is probed at time n and $\mathbf{1}_k(n) = 0$ otherwise. Let t_m be the time when the decision maker starts the m^{th} test. For example, assume that $K = 3$ and the decision maker tests the components according to the following order: 3, 1, 2. Then, $t_1 = 1$ (when the test starts), $t_2 = \tau_3 + 1$, $t_3 = \tau_1 + 1$.

Under the independent model, the posterior probability of component k being abnormal can be updated at time t_{m+1} as follows [22]:

$$\begin{aligned} \pi_k(t_{m+1}) &= (1 - \mathbf{1}_k(t_m)) \pi_k(t_m) \\ &+ \frac{\mathbf{1}_k(t_m) \pi_k(t_m) f_k^{(1)}(\mathbf{y}_k(N_k))}{\pi_k(t_m) f_k^{(1)}(\mathbf{y}_k(N_k)) + (1 - \pi_k(t_m)) f_k^{(0)}(\mathbf{y}_k(N_k))}, \end{aligned} \quad (25)$$

where $\pi_k(t_1) = \pi_k$ denotes the *a priori* probability of component k being abnormal. The term $\mathbf{y}_k(N_k) = \{y_k(i)\}_{i=t_m}^{t_m+N_k-1}$ denotes the N_k -size vector of observations, taken from component k . Under the exclusive model, $\pi_k(t_{m+1})$ is given in (26) at the top of the next page. Note that in contrast to the independent model, under the exclusive model the beliefs of all the components are changed at each time due to the dependency across components. The posterior probabilities depend on the selection rule and the collected measurements.

$$\pi_k(t_{m+1}) = \frac{\mathbf{1}_k(t_m)\pi_k(t_m)f_k^{(1)}(\mathbf{y}_k(N_k))}{\pi_k(t_m)f_k^{(1)}(\mathbf{y}_k(N_k)) + (1 - \pi_k(t_m))f_k^{(0)}(\mathbf{y}_k(N_k))} + \frac{(1 - \mathbf{1}_k(t_m))\pi_k(t_m)f_{\phi(t_m)}^{(0)}(\mathbf{y}_{\phi(t_m)}(N_{\phi(t_m)}))}{\pi_{\phi(t_m)}(t_m)f_{\phi(t_m)}^{(1)}(\mathbf{y}_{\phi(t_m)}(N_{\phi(t_m)})) + (1 - \pi_{\phi(t_m)}(t_m))f_{\phi(t_m)}^{(0)}(\mathbf{y}_{\phi(t_m)}(N_{\phi(t_m)}))}. \quad (26)$$

A. Proof of Theorem 1 Under The Exclusive Model

Let $\mathbf{E}'(N_k|H_i, t)$ be the expected sample size achieved by a stopping rule and a decision rule $(\tau'_k(t), \delta'_k(t))$, depending on the time that component k is tested (i.e., $(\tau'_k(t), \delta'_k(t))$ depend on the selection rule), such that error constraints are satisfied. Let $\mathbf{E}^{A2}(N_k|H_i)$ be the expected sample size achieved by the SPRT's stopping rule and decision rule $(\tau_k^{A2}, \delta_k^{A2})$, independent of the time that component k is tested (i.e., $(\tau_k^{A2}, \delta_k^{A2})$ are independent of the selection rule), such that error constraints are satisfied. Clearly, $\mathbf{E}^{A2}(N_k|H_i) \leq \mathbf{E}'(N_k|H_i, t)$ for all k, t , for $i = 0, 1$.

Step 1: Proving the theorem for $K = 2$:

Assume that

$$\frac{\pi_1(t_1)c_1}{\mathbf{E}^{A2}(N_1|H_0)} \geq \frac{\pi_2(t_1)c_2}{\mathbf{E}^{A2}(N_2|H_0)}. \quad (27)$$

Consider selection rules $\phi^{(1)}, \phi^{(2)}$ that select component 1 first followed by component 2 and component 2 first followed by component 1, respectively. The expected cost achieved by $(\tau'(t), \delta'(t), \phi^{(2)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau'(t), \delta'(t), \phi^{(2)}) \right\} \\ &= (\mathbf{E}'(N_2|H_1, t_1)) \pi_2(t_1)c_2 \\ &+ (\mathbf{E}'(N_2|H_0, t_1) + \mathbf{E}'(N_1|H_1, t_2)) \pi_1(t_1)c_1. \end{aligned} \quad (28)$$

The expected cost achieved by $(\tau'(t), \delta'(t), \phi^{(1)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau'(t), \delta'(t), \phi^{(1)}) \right\} \\ &= (\mathbf{E}'(N_1|H_1, t_1)) \pi_1(t_1)c_1 \\ &+ (\mathbf{E}'(N_1|H_0, t_1) + \mathbf{E}'(N_2|H_1, t_2)) \pi_2(t_1)c_2. \end{aligned} \quad (29)$$

Note that the expected cost achieved by both selection rules can be further reduced by minimizing the expected sample sizes (such that error constraints are satisfied) independent of the selection rules, which

is achieved by $(\tau_k^{A2}, \delta_k^{A2})$. Therefore, an optimal solution must be $(\tau^{A2}, \delta^{A2}, \phi^{(1)})$ or $(\tau^{A2}, \delta^{A2}, \phi^{(2)})$. Next, we use the interchange argument to prove the theorem for $K = 2$. The expected cost achieved by $(\tau^{A2}, \delta^{A2}, \phi^{(2)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A2}, \delta^{A2}, \phi^{(2)}) \right\} \\ &= (\mathbf{E}^{A2}(N_2|H_1)) \pi_2(t_1) c_2 \\ & \quad + (\mathbf{E}^{A2}(N_2|H_0) + \mathbf{E}^{A2}(N_1|H_1)) \pi_1(t_1) c_1. \end{aligned} \quad (30)$$

The expected cost achieved by $(\tau^{A2}, \delta^{A2}, \phi^{(1)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A2}, \delta^{A2}, \phi^{(1)}) \right\} \\ &= (\mathbf{E}^{A2}(N_1|H_1)) \pi_1(t_1) c_1 \\ & \quad + (\mathbf{E}^{A2}(N_1|H_0) + \mathbf{E}^{A2}(N_2|H_1)) \pi_2(t_1) c_2. \end{aligned} \quad (31)$$

the expected cost achieved by $\phi^{(1)}$ is lower than that achieved by $\phi^{(2)}$ since $\frac{\pi_1(t_1)c_1}{\mathbf{E}^{A2}(N_1|H_0)} \geq \frac{\pi_2(t_1)c_2}{\mathbf{E}^{A2}(N_2|H_0)}$, which completes the proof for $K = 2$.

Step 2: Proving the theorem by induction on the number of components K :

Assume that the theorem is true for $K - 1$ components (where one and only one component is abnormal).

Assume that

$$\frac{\pi_1(t_1)c_1}{\mathbf{E}^{A2}(N_1|H_0)} \geq \frac{\pi_2(t_1)c_2}{\mathbf{E}^{A2}(N_2|H_0)} \geq \dots \geq \frac{\pi_K(t_1)c_K}{\mathbf{E}^{A2}(N_K|H_0)}. \quad (32)$$

Consider the case of K components and denote $\phi^{(j)}$ as an optimal selection rule that selects component j first.

Step 2.1: Proving the theorem for the last $K - 1$ components:

Next, we show that the last $K - 1$ components must be selected in decreasing order of $\pi_k(t_1)c_k/\mathbf{E}^{A2}(N_k|H_0)$ and tested by the SPRT.

Let

$$\gamma_j(t) = \frac{1}{\pi_j(t) \frac{f_j^{(1)}(\mathbf{y}_j(N_j))}{f_j^{(0)}(\mathbf{y}_j(N_j))} + 1 - \pi_j(t)}. \quad (33)$$

Note that when the decision maker completes testing component j , the other components update their beliefs according to:

$$\pi_k(t_2) = \gamma_j(t_1) \pi_k(t_1), \quad \forall k \neq j. \quad (34)$$

The expected cost achieved by $\phi^{(j)}$ given the outcome (at time t_2) by testing component j (i.e., given the observations vector $\mathbf{y}_j(N_j)$) is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid \phi^{(j)}, \mathbf{y}_j(N_j) \right\} \\ &= \pi_j(t_2) c_j N_j + (1 - \pi_j(t_2)) \times \\ & \mathbf{E} \left\{ \sum_{k=1, k \neq j}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid \phi^{(j)}, \mathbf{y}_j(N_j), j \in \mathcal{H}_0 \right\}. \end{aligned} \quad (35)$$

Let

$$\tilde{\tau}_k = \tau_k - N_j \quad \forall k \neq j \quad (36)$$

be the modified stopping time, defined as the stopping time from $t = N_j + 1$ until testing of component k is completed. Thus, we can rewrite (35) as:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid \phi^{(j)}, \mathbf{y}_j(N_j) \right\} \\ &= \sum_{k=1}^K \pi_k(t_2) c_k N_j + (1 - \pi_j(t_2)) \times \\ & \mathbf{E} \left\{ \sum_{k=1, k \neq j}^K c_k \tilde{\tau}_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid \phi^{(j)}, \mathbf{y}_j(N_j), j \in \mathcal{H}_0 \right\}. \end{aligned} \quad (37)$$

The term $\sum_{k=1}^K \pi_k(t_2) c_k N_j$ in (37) follows since,

$$\begin{aligned} & \Pr(k \in \mathcal{H}_1 \mid \phi^{(j)}, \mathbf{y}_j(N_j), j \in \mathcal{H}_0) \\ &= \frac{\Pr(k \in \mathcal{H}_1, j \in \mathcal{H}_0 \mid \phi^{(j)}, \mathbf{y}_j(N_j))}{\Pr(j \in \mathcal{H}_0 \mid \phi^{(j)}, \mathbf{y}_j(N_j))} \\ &= \frac{\Pr(k \in \mathcal{H}_1 \mid \phi^{(j)}, \mathbf{y}_j(N_j))}{\Pr(j \in \mathcal{H}_0 \mid \phi^{(j)}, \mathbf{y}_j(N_j))} = \frac{\pi_k(t_2)}{1 - \pi_j(t_2)} \triangleq \tilde{\pi}_k(t_2). \end{aligned} \quad (38)$$

Minimizing

$$\mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid \phi^{(j)}, \mathbf{y}_j(N_j) \right\} \quad (39)$$

at time t_2 , requires one to minimize

$$\mathbf{E} \left\{ \sum_{k=1, k \neq j}^K c_k \tilde{\tau}_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid \phi^{(j)}, \mathbf{y}_j(N_j), j \in \mathcal{H}_0 \right\} \quad (40)$$

in (37).

Note that (40) is the cost for $K - 1$ components (where one and only one component is abnormal) starting

at time $t = t_2 = N_j + 1$, with prior probability $\tilde{\pi}_k(t_2) = \frac{\pi_k(t_2)}{1 - \pi_j(t_2)}$ for component $k \neq j$ being abnormal. By the induction hypothesis, for any optimal selection rule $\phi^{(j)}$ that selects component j first, arranging the last $K - 1$ components in decreasing order of $\tilde{\pi}_k(t_2)c_k/\mathbf{E}^{A2}(N_k|H_0)$ (and testing them by the SPRT) minimizes (40).

Since

$$\tilde{\pi}_k(t_2) = \frac{\gamma_j(t_1)}{1 - \pi_j(t_2)} \pi_k(t_1) \quad \forall k \neq j, \quad (41)$$

then

$$\begin{aligned} \frac{\tilde{\pi}_1(t_2)c_1}{\mathbf{E}^{A2}(N_1|H_0)} &\geq \frac{\tilde{\pi}_2(t_2)c_2}{\mathbf{E}^{A2}(N_2|H_0)} \geq \dots \geq \frac{\tilde{\pi}_{j-1}(t_2)c_{j-1}}{\mathbf{E}^{A2}(N_{j-1}|H_0)} \\ &\geq \frac{\tilde{\pi}_{j+1}(t_2)c_{j+1}}{\mathbf{E}^{A2}(N_{j+1}|H_0)} \geq \dots \geq \frac{\tilde{\pi}_K(t_2)c_K}{\mathbf{E}^{A2}(N_K|H_0)}. \end{aligned} \quad (42)$$

Thus, the last $K - 1$ components must be selected in decreasing order of $\pi_k(t_1)c_k/\mathbf{E}^{A2}(N_k|H_0)$ and tested by the SPRT.

Step 2.2: Proving the theorem for all the K components:

Finally, we show that component 1 (i.e., the component with the highest index) must be selected first.

The expected cost achieved by $(\tau'(t), \delta'(t), \phi^{(j)})$ is given by:

$$\begin{aligned} &\mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau'(t), \delta'(t), \phi^{(j)}) \right\} \\ &= \pi_j(t_1)c_j (\mathbf{E}'(N_j|H_1, t_1)) + \sum_{k=1, k \neq j}^K [\pi_k(t_1)c_k \times \\ &\quad \left(\mathbf{E}'(N_j|H_0, t_1) + \left(\sum_{i=1, i \neq j}^{k-1} \mathbf{E}^{A2}(N_i|H_0) \right) \right. \\ &\quad \left. + \mathbf{E}^{A2}(N_k|H_1) \right)] . \end{aligned} \quad (43)$$

First, note that the expected cost achieved by $(\tau'(t), \delta'(t), \phi^{(j)})$ can be further reduced for all j by minimizing the expected sample size $\mathbf{E}'(N_j|H_i, t_1)$ for $i = 0, 1$, which is achieved by $(\tau_j^{A2}, \delta_j^{A2})$. Therefore, an optimal solution must be $(\tau^{A2}, \delta^{A2}, \phi^{(j)})$ for an optimal selection rule $\phi^{(j)}$. Thus, in the following we consider solutions of the form $(\tau^{A2}, \delta^{A2}, \phi)$.

Next, by contradiction, consider an optimal selection rule $\phi^{(j \neq 1)}$ that selects component $j \neq 1$ first. Therefore, $\phi^{(j \neq 1)}$ selects the components in the following order:

$$j, 1, 2, \dots, j-1, j+1, \dots, K.$$

As a result, the expected cost achieved by $(\tau^{A2}, \delta^{A2}, \phi^{(j \neq 1)})$ is given by:

$$\begin{aligned}
& \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A2}, \delta^{A2}, \phi^{(j \neq 1)}) \right\} \\
&= \pi_j(t_1) c_j (\mathbf{E}^{A2}(N_j | H_1)) \\
&\quad + \pi_1(t_1) c_1 [\mathbf{E}^{A2}(N_j | H_0) + \mathbf{E}^{A2}(N_1 | H_1)] \\
&\quad + \sum_{k=2, k \neq j}^K [\pi_k(t_1) c_k \times \\
&\quad \left(\mathbf{E}^{A2}(N_j | H_0) + \left(\sum_{i=1, i \neq j}^{k-1} \mathbf{E}^{A2}(N_i | H_0) \right) \right. \\
&\quad \left. + \mathbf{E}^{A2}(N_k | H_1) \right)] .
\end{aligned} \tag{44}$$

We use the interchange argument to prove the theorem. Consider a selection rule $\phi^{(1)}$ that selects component 1 first followed by components $j, 2, 3, j-1, j+1, \dots, K$. Similar to (44), the expected cost achieved by $(\tau^{A2}, \delta^{A2}, \phi^{(1)})$ is given by:

$$\begin{aligned}
& \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A2}, \delta^{A2}, \phi^{(1)}) \right\} \\
&= \pi_1(t_1) c_1 (\mathbf{E}^{A2}(N_1 | H_1)) \\
&\quad + \pi_j(t_1) c_j [\mathbf{E}^{A2}(N_1 | H_0) + \mathbf{E}^{A2}(N_j | H_1)] \\
&\quad + \sum_{k=2, k \neq j}^K [\pi_k(t_1) c_k \times \\
&\quad \left(\mathbf{E}^{A2}(N_j | H_0) + \left(\sum_{i=1, i \neq j}^{k-1} \mathbf{E}^{A2}(N_i | H_0) \right) \right. \\
&\quad \left. + \mathbf{E}^{A2}(N_k | H_1) \right)] .
\end{aligned} \tag{45}$$

By comparing (44) and (45), it can be verified that:

$$\begin{aligned}
& \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A2}, \delta^{A2}, \phi^{(1)}) \right\} \\
&\leq \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A2}, \delta^{A2}, \phi^{(j \neq 1)}) \right\}
\end{aligned}$$

since $\pi_1(t_1) c_1 / \mathbf{E}^{A2}(N_1 | H_0) \geq \pi_j(t_1) c_j / \mathbf{E}^{A2}(N_j | H_0)$.

The expected cost can be reduced by selecting component 1 first followed by component j , which contradicts the optimality of $\phi^{(j \neq 1)}$. Hence, at time t_1 selecting component 1 minimizes the expected cost. We

have already proved that selecting the last $K-1$ components in decreasing order of $\pi_k(t_1)c_k/\mathbf{E}^{A2}(N_k|H_0)$ minimizes the objective function, which completes the proof. ■

B. Proof of Theorem 1 Under The Independent Model

Let $\mathbf{E}'(N_k|H_i, t)$ be the expected sample size achieved by a stopping rule and a decision rule $(\tau'_k(t), \delta'_k(t))$, depending on the time that component k is tested (i.e., $(\tau'_k(t), \delta'_k(t))$ depend on the selection rule), such that error constraints are satisfied. Let $\mathbf{E}^{A1}(N_k|H_i)$ be the expected sample size achieved by the SPRT's stopping rule and decision rule $(\tau_k^{A1}, \delta_k^{A1})$, independent of the time that component k is tested (i.e., $(\tau_k^{A1}, \delta_k^{A1})$ are independent of the selection rule), such that error constraints are satisfied. Clearly, $\mathbf{E}^{A1}(N_k|H_i) \leq \mathbf{E}'(N_k|H_i, t)$ for all k, t , for $i = 0, 1$ and are achieved by the πcN -SPRT algorithm.

First, consider the case where $K = 2$. Assume that

$$\frac{\pi_1(t_1)c_1}{\mathbf{E}^{A1}(N_1)} \geq \frac{\pi_2(t_1)c_2}{\mathbf{E}^{A1}(N_2)}.$$

Consider selection rules $\phi^{(1)}, \phi^{(2)}$ that select component 1 first followed by component 2 and component 2 first followed by component 1, respectively. The expected cost achieved by $(\tau'(t), \delta'(t), \phi^{(2)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau'(t), \delta'(t), \phi^{(2)}) \right\} \\ &= (\mathbf{E}'(N_2|H_1, t_1)) \pi_2(t_1)c_2 \\ &+ (\mathbf{E}'(N_2|t_1) + \mathbf{E}'(N_1|H_1, t_2)) \pi_1(t_1)c_1. \end{aligned} \tag{46}$$

The expected cost achieved by $(\tau'(t), \delta'(t), \phi^{(1)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau'(t), \delta'(t), \phi^{(1)}) \right\} \\ &= (\mathbf{E}'(N_1|H_1, t_1)) \pi_1(t_1)c_1 \\ &+ (\mathbf{E}'(N_1|t_1) + \mathbf{E}'(N_2|H_1, t_2)) \pi_2(t_1)c_2. \end{aligned} \tag{47}$$

Note that the expected cost achieved by both selection rules can be further reduced by minimizing the expected sample sizes (such that error constraints are satisfied) independent of the selection rules, which is achieved by $(\tau_k^{A1}, \delta_k^{A1})$. Therefore, an optimal solution must be $(\tau^{A1}, \delta^{A1}, \phi^{(1)})$ or $(\tau^{A1}, \delta^{A1}, \phi^{(2)})$. Next, we use the interchange argument to prove the theorem for $K = 2$. The expected cost achieved by

$(\tau^{A1}, \delta^{A1}, \phi^{(2)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A1}, \delta^{A1}, \phi^{(2)}) \right\} \\ &= (\mathbf{E}^{A1}(N_2 | H_1)) \pi_2(t_1) c_2 \\ &+ (\mathbf{E}^{A1}(N_2) + \mathbf{E}^{A1}(N_1 | H_1)) \pi_1(t_1) c_1. \end{aligned} \quad (48)$$

The expected cost achieved by $(\tau^{A1}, \delta^{A1}, \phi^{(1)})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A1}, \delta^{A1}, \phi^{(1)}) \right\} \\ &= (\mathbf{E}^{A1}(N_1 | H_1)) \pi_1(t_1) c_1 \\ &+ (\mathbf{E}^{A1}(N_1) + \mathbf{E}^{A1}(N_2 | H_1)) \pi_2(t_1) c_2. \end{aligned} \quad (49)$$

The expected cost achieved by $\phi^{(1)}$ is lower than that achieved by $\phi^{(2)}$ since $\frac{\pi_1(t_1)c_1}{\mathbf{E}^{A1}(N_1)} \geq \frac{\pi_2(t_1)c_2}{\mathbf{E}^{A1}(N_2)}$, which completes the proof for $K = 2$.

The rest of the proof follows by induction on the number of components, as was done under the exclusive model. ■

C. Proof of Theorem 2

For every k , let $\mathbf{E}^*(N_k | H_i)$ be the minimal expected sample size that can be achieved by any sequential test, such that error constraints are satisfied. Let $\mathbf{E}^{A3}(N_k | H_i)$ be the expected sample size achieved by the πcN -SGLRT/SALRT algorithm, such that error constraints are satisfied. Clearly, $\mathbf{E}^*(N_k | H_i) \leq \mathbf{E}^{A3}(N_k | H_i)$ for all k , for $i = 0, 1$.

Assume that

$$\frac{\pi_1(t_1)c_1}{\mathbf{E}^*(N_1)} \geq \frac{\pi_2(t_1)c_2}{\mathbf{E}^*(N_2)} \geq \dots \geq \frac{\pi_K(t_1)c_K}{\mathbf{E}^*(N_K)}. \quad (50)$$

Similar to the proof of Theorem 1, it can be verified that the optimal solution to (2) is to select the components in the following order: $1, 2, \dots, K$, where the components are tested by a sequential test that achieves expected sample size $\mathbf{E}^*(N_k | H_i)$ for all k , for $i = 0, 1$. Therefore, the expected cost achieved by $(\tau^*, \delta^*, \phi^*)$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^*, \delta^*, \phi^*) \right\} \\ &= \sum_{k=1}^K \pi_k(t_1) c_k \left[\left(\sum_{i=1}^{k-1} \mathbf{E}^*(N_i) \right) + \mathbf{E}^*(N_k | H_1) \right]. \end{aligned} \quad (51)$$

By the asymptotic optimality property of the SALRT/SGLRT for a single process (used in the πcN -SGLRT/SALRT algorithm), it follows that $\mathbf{E}^{A3}(N_k|H_i) \sim \mathbf{E}^*(N_k|H_i)$ for all k , for $i = 0, 1$ as $P_k^{FA} \rightarrow 0$, $P_k^{MD} \rightarrow 0$. As a result, for sufficiently small error probabilities, the solution $(\tau^{A3}, \delta^{A3}, \phi^{A3})$ is to select the components in the following order: $1, 2, \dots, K$, where the components are tested by an asymptotically optimal sequential test that achieves expected sample size $\mathbf{E}^{A3}(N_k|H_i)$ for all k , for $i = 0, 1$. Therefore, the expected cost achieved by $(\tau^{A3}, \delta^{A3}, \phi^{A3})$ is given by:

$$\begin{aligned} & \mathbf{E} \left\{ \sum_{k=1}^K c_k \tau_k \mathbf{1}_{\{k \in \mathcal{H}_1\}} \mid (\tau^{A3}, \delta^{A3}, \phi^{A3}) \right\} \\ &= \sum_{k=1}^K \pi_k(t_1) c_k \left[\left(\sum_{i=1}^{k-1} \mathbf{E}^{A3}(N_i) \right) + \mathbf{E}^{A3}(N_k|H_1) \right]. \end{aligned} \quad (52)$$

Since $\mathbf{E}^{A3}(N_k|H_i) \sim \mathbf{E}^*(N_k|H_i)$ for $i = 0, 1$ as $P_k^{FA} \rightarrow 0$, $P_k^{MD} \rightarrow 0$ for all k , the theorem follows. ■

D. Proof of Theorem 3

The structure of the proof is similar to the proof of Theorem 2. Hence, we provide a sketch of the proof, using notation similar to that used in the proof of Theorem 2. Similar to the proof of Theorem 1, it can be verified that the optimal solution to (2) is to select the components in decreasing order of $\pi_k(t_1)c_k/\mathbf{E}^*(N_k|H_0)$, where the components are tested by a sequential test that achieves expected sample size $\mathbf{E}^*(N_k|H_i)$ for all k , for $i = 0, 1$. By the asymptotic optimality property for a single process of the SALRT/SGLRT (used in the πcN_0 -SGLRT/SALRT algorithm), it follows that $\mathbf{E}^{A4}(N_k|H_i) \sim \mathbf{E}^*(N_k|H_i)$ for all k , for $i = 0, 1$ as $P_k^{FA} \rightarrow 0$, $P_k^{MD} \rightarrow 0$. As a result, for sufficiently small error probabilities, the solution $(\tau^{A4}, \delta^{A4}, \phi^{A4})$ is to select the components in decreasing order of $\pi_k(t_1)c_k/\mathbf{E}^*(N_k|H_0)$, where the components are tested by an asymptotically optimal sequential test that achieves expected sample size $\mathbf{E}^{A4}(N_k|H_i)$ for all k , for $i = 0, 1$. Similar to the proof of Theorem 2, comparing the objective functions achieved by $(\tau^*, \delta^*, \phi^*)$ and $(\tau^{A4}, \delta^{A4}, \phi^{A4})$ proves the theorem. ■

REFERENCES

- [1] Q. Zhao and B. M. Sadler, "A survey of dynamic spectrum access," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 79–89, 2007.
- [2] A. Wald, "Sequential analysis," *New York: Wiley*, 1947.
- [3] G. Schwarz, "Asymptotic shapes of Bayes sequential testing regions," *The Annals of mathematical statistics*, pp. 224–236, 1962.
- [4] H. Robbins and D. Siegmund, "The expected sample size of some tests of power one," *The Annals of Statistics*, pp. 415–436, 1974.

- [5] T. L. Lai, "Nearly optimal sequential tests of composite hypotheses," *The Annals of Statistics*, pp. 856–886, 1988.
- [6] I. V. Pavlov, "Sequential procedure of testing composite hypotheses with applications to the Kiefer-Weiss problem," *Theory of Probability and Its Applications*, vol. 35, no. 2, pp. 280–292, 1990.
- [7] T. L. Lai and L. M. Zhang, "Nearly optimal generalized sequential likelihood ratio tests in multivariate exponential families," *Lecture Notes-Monograph Series*, pp. 331–346, 1994.
- [8] A. G. Tartakovsky, "An efficient adaptive sequential procedure for detecting targets," in *IEEE Aerospace Conference Proceedings, 2002*, vol. 4, pp. 1581–1596, 2002.
- [9] V. Draglin, A. G. Tartakovsky, and V. V. Veeravalli, "Multihypothesis sequential probability ratio tests - part i: Asymptotic optimality," *IEEE Transactions on Information Theory*, vol. 45, no. 7, pp. 2448–2461, 1999.
- [10] Q. Zhao and J. Ye, "Quickest detection in multiple on-off processes," *IEEE Transactions on Signal Processing*, vol. 58, no. 12, pp. 5994–6006, 2010.
- [11] H. Li, "Restless watchdog: Selective quickest spectrum sensing in multichannel cognitive radio systems," *EURASIP Journal on Advances in Signal Processing*, 2009.
- [12] R. Caromi, Y. Xin, and L. Lai, "Fast multiband spectrum scanning for cognitive radio systems," *IEEE Transaction on Communications*, vol. 61, no. 1, pp. 63–75, 2013.
- [13] L. Lai, H. V. Poor, Y. Xin, and G. Georgiadis, "Quickest search over multiple sequences," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 5375–5386, 2011.
- [14] M. L. Malloy, G. Tang, and R. D. Nowak, "Quickest search for a rare distribution," *IEEE Annual Conference on Information Sciences and Systems*, pp. 1–6, 2012.
- [15] A. Tajer and H. V. Poor, "Quick search for rare events," *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4462–4481, 2013.
- [16] K. S. Zigangirov, "On a problem in optimal scanning," *Theory of Probability and Its Applications*, vol. 11, no. 2, pp. 294–298, 1966.
- [17] E. Klimko and J. Yackel, "Optimal search strategies for Wiener processes," *Stochastic Processes and their Applications*, vol. 3, no. 1, pp. 19–33, 1975.
- [18] V. Dragalin, "A simple and effective scanning rule for a multi-channel system," *Metrika*, vol. 43, no. 1, pp. 165–182, 1996.
- [19] L. D. Stone and J. A. Stanshine, "Optimal search using uninterrupted contact investigation," *SIAM Journal on Applied Mathematics*, vol. 20, no. 2, pp. 241–263, 1971.
- [20] K. P. Tognetti, "An optimal strategy for a whereabouts search," *Operations Research*, vol. 16, no. 1, pp. 209–211, 1968.
- [21] J. B. Kadane, "Optimal whereabouts search," *Operations Research*, vol. 19, no. 4, pp. 894–904, 1971.
- [22] D. A. Castanon, "Optimal search strategies in dynamic hypothesis testing," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 25, no. 7, pp. 1130–1138, 1995.
- [23] Y. Zhai and Q. Zhao, "Dynamic search under false alarms," in *Proc. of the IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2013.
- [24] D. Blackwell, "Equivalent comparisons of experiments," *The Annals of Mathematical Statistics*, vol. 24, no. 2, pp. 265–272, 1953.
- [25] H. Chernoff, "Sequential design of experiments," *The Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 755–770, 1959.
- [26] M. H. DeGroot, "Uncertainty, information, and sequential experiments," *The Annals of Mathematical Statistics*, pp. 404–419, 1962.

- [27] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *The Annals of Statistics*, vol. 41, no. 6, pp. 2703–2738, 2013.
- [28] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2451–2464, 2013.
- [29] K. Cohen and Q. Zhao, "Active hypothesis testing for quickest anomaly detection," *submitted to IEEE Trans. on Information Theory*, 2014. Available in *arXiv:1403.1023v1 [cs.IT]* at <http://arxiv.org/abs/1403.1023>.
- [30] K. Cohen and Q. Zhao, "Quickest anomaly detection: a case of active hypothesis testing," in *Proc. of Information Theory and Applications (ITA) Workshop*, 2014.
- [31] W. E. Smith, "Various optimizers for single-stage production," *Naval Research Logistics Quarterly*, vol. 3, no. 1-2, pp. 59–66, 1956.
- [32] I. Onat and A. Miri, "An intrusion detection system for wireless sensor networks," in *IEEE International Conference on Wireless And Mobile Computing, Networking And Communications*, vol. 3, pp. 253–259, 2005.